

The Patent Office
Concept House
Cardiff Road
Newport
South Wales
NP10 8QQ

J1036 U.S. PTO
09/858437



I, the undersigned, being an officer duly authorised in accordance with Section 74(1) and (4) of the Deregulation & Contracting Out Act 1994, to sign and issue certificates on behalf of the Comptroller-General, hereby certify that annexed hereto is a true copy of the documents as originally filed in connection with the patent application identified therein.

In accordance with the Patents (Companies Re-registration) Rules 1982, if a company named in this certificate and any accompanying documents has re-registered under the Companies Act 1980 with the same name as that with which it was registered immediately before re-registration save for the substitution as, or inclusion as, the last part of the name of the words "public limited company" or their equivalents in Welsh, references to the name of the company in this certificate and any accompanying documents shall be treated as references to the name with which it is so re-registered.

In accordance with the rules, the words "public limited company" may be replaced by p.l.c., plc, P.L.C. or PLC.


Re-registration under the Companies Act does not constitute a new legal entity but merely subjects the company to certain additional company law rules.

Signed

Dated

20 March 2001

THIS PAGE BLANK (USPTO)



1/77

25MAY00 E539584-1 D00107
P01/7700 0.00-0012592 The Patent Office

Cardiff Road
Newport
South Wales
NP10 8QQ

Request for grant of a patent

(See the notes on the back of this form. You can also get an explanatory leaflet from the Patent Office to help you fill in this form)

1. Your reference

N00/0374/GB

2. Patent application number

(The Patent Office will fill in this part)

24 MAY 2000

0012592.2

3. Full name, address and postcode of the or of each applicant (underline all surnames)

Power X Limited
Stafford Court
145 Washway Road
Sale
Cheshire M33 7PE

Patents ADP number (if you know it)

If the applicant is a corporate body, give the country/state of its incorporation

6803233002

4. Title of the invention

High Speed Digital Switch Arbiter

5. Name of your agent (if you have one)

"Address for service" in the United Kingdom to which all correspondence should be sent (including the postcode)

McNeight & Lawrence
Regent House, Heaton Lane
Stockport, Cheshire SK4 1BS

Patents ADP number (if you know it)

0001115001

6. If you are declaring priority from one or more earlier patent applications, give the country and the date of filing of the or of each of these earlier applications and (if you know it) the or each application number

Country

Priority application number
(if you know it)

Date of filing
(day / month / year)

7. If this application is divided or otherwise derived from an earlier UK application, give the number and the filing date of the earlier application

Number of earlier application

Date of filing
(day / month / year)

8. Is a statement of inventorship and of right to grant of a patent required in support of this request? (Answer 'Yes' if:

Yes

- a) any applicant named in part 3 is not an inventor, or
 - b) there is an inventor who is not named as an applicant, or
 - c) any named applicant is a corporate body.
- See note (d))

Patents Form 1/77

9. Enter the number of sheets for any of the following items you are filing with this form. Do not count copies of the same document

Continuation sheets of this form

Description 4

Claim(s)

Abstract

Drawing(s) 3

+3 12

10. If you are also filing any of the following, state how many against each item.

Priority documents

Translations of priority documents

Statement of inventorship and right to grant of a patent (*Patents Form 7/77*)

Request for preliminary examination and search (*Patents Form 9/77*)

Request for substantive examination (*Patents Form 10/77*)

Any other documents
(please specify)

11. I/We request the grant of a patent on the basis of this application.

Signature

Date

McNeighr & Lawrence

23 May 2000

12. Name and daytime telephone number of person to contact in the United Kingdom
- J G Lawrence
0161 480 6394

Warning

After an application for a patent has been filed, the Comptroller of the Patent Office will consider whether publication or communication of the invention should be prohibited or restricted under Section 22 of the Patents Act 1977. You will be informed if it is necessary to prohibit or restrict your invention in this way. Furthermore, if you live in the United Kingdom, Section 23 of the Patents Act 1977 stops you from applying for a patent abroad without first getting written permission from the Patent Office unless an application has been filed at least 6 weeks beforehand in the United Kingdom for a patent for the same invention and either no direction prohibiting publication or communication has been given, or any such direction has been revoked.

Notes

- a) If you need help to fill in this form or you have any questions, please contact the Patent Office on 0645 500505.
- b) Write your answers in capital letters using black ink or you may type them.
- c) If there is not enough space for all the relevant details on any part of this form, please continue on a separate sheet of paper and write "see continuation sheet" in the relevant part(s). Any continuation sheet should be attached to this form.
- d) If you have answered 'Yes' Patents Form 7/77 will need to be filed.
- e) Once you have filled in the form you must remember to sign and date it.
- f) For details of the fee and ways to pay please contact the Patent Office.

Title: High speed digital switch arbiter

The present invention relates to arbiters for use in data switching systems and is more particularly concerned with so-called "bipartite graph" matching for use in arbitrating between connection requests for cross bar connections in a data switching system.

At present, bipartite graph matching algorithms that can be implemented in hardware are based upon a Round Robin pointer manipulation schemes. The purpose of the round robin pointers is to find the location of the next valid value out of a finite set of values. If no valid values are found the location of the pointer does not move. When a three-stage arbiter is constructed from just round robin pointers, under high load, efficiency is reduced due to an effect known as pointer synchronisation.

Several methods attempt to solve the pointer synchronisation problem, one of which is SLIP (McKeown US005500858A).

Round Robin Arbiters

In computing and networking, there is often contention for a specific resource, e.g. processor time or bandwidth. One of the simplest ways to allocate resources it to grant access to each requester in turn, this is generally known as a round robin.

When a round robin is implemented, some form of pointer is used to indicate the last process that was serviced. The simplest implementation of the round robin moves from one location to the next in its sequence once every operation cycle. If there is nothing at the location selected then nothing is done until the next cycle. A logical extension to the Round Robin allows the pointer to skip locations that do not have any requests for the resource; this increases the efficiency of the algorithm.

Three Stage Arbiters

As mentioned above Round Robin arbiters are used when there are many processes requiring access to one resource. The Round Robin is a form of two stage arbiter i.e. requests are presented and then one is selected. In the case of a crossbar switch arbiter there are many services requesting access to many resources, otherwise known as a Bipartite Graph Matching problem. An extra final stage is added to the operation to allow for multiple resources.

The first two stages of the arbitration are the same as a Round Robin requests are presented to the resources and each grants one. However, as there are multiple resources each of the requesting processes must also select one from many in order to ensure exclusive access.

Pointer Synchronisation

When a round robin arbiter is placed under full load the position of the grant pointers do not change relative to one another. If multiple grant pointers are pointing to any accept stage only one can be accepted, causing a reduction in the overall efficiency of the arbiter and an increase in the effective load. Any inefficiency that arises in the arbiter will therefore be self-sustaining. Some method of changing the pointers relative position is therefore required.

It is an object of the present invention to overcome the above mentioned pointer synchronisation problem.

According to the invention there is provided an arbitration arrangement for use in selecting the connections to be made between ingress and egress ports of a memoryless cross-bar switch of a data switching system, the arbitration arrangement comprising a three phase process involving (i) a request phase in which each ingress port sends its connection requests to the egress ports to which a connection is required, (ii) a grant phase in which each egress port examines in a round-robin manner the requests directed to it and selects one request for grant returning a grant signal indicative of the selected request and directed to the appropriate ingress port and (iii) an accept phase in which each ingress port examines in a round-robin manner the received grant signals and selects one to accept thereby defining an ingress to egress port connection across the cross-bar switch characterised in that the request and grant signal paths are connected in a matrix so as to ensure that the grant round-robin transition sequences are mutually exclusive for each egress port.

According to a feature of the invention there is provided an arbitration arrangement in which the connection matrix conforms to a formula of

$$M^* = \left(M + \frac{P(P+1)}{2} \right) \bmod N$$

where M^* = mapped pointer location; M = round-robin pointer location;

P = port number and N = number of cross-bar switch ports.

Also according to the invention there is provided an arbitration process for use in selecting the connections to be made between ingress and egress ports of a memoryless

cross-bar switch of a data switching system, the arbitration arrangement comprising a three phase process involving (i) a request phase in which each ingress port sends its connection requests to the egress ports to which a connection is required, (ii) a grant phase in which each egress port examines in a round-robin manner the requests directed to it and selects one request for grant returning a grant signal indicative of the selected request and directed to the appropriate ingress port and (iii) an accept phase in which each ingress port examines in a round-robin manner the received grant signals and selects one to accept thereby defining an ingress to egress port connection across the cross-bar switch characterised in that the request and grant signal paths are connected in a matrix so as to ensure that the grant round-robin transition sequences are mutually exclusive for each egress port.

The worst case for pointer synchronisation is all of the grants going to one location; an arbiter in this state will have its efficiency reduced to 1/number of ports. The ideal solution for this would be to have all the pointers move to different locations on the next arbitration cycle. One-way of achieving this is to give every round robin a different path through its set of requests.

The simplest way to cause the grant round robin pointers to appear to follow different paths is to reconnect the request and grant matrix signals in a new order (see Figure 1). A useful aspect of reordering the pointers in this way is that no extra gates are required to implement the required changes.

In order to totally avoid the problem of synchronisation the path for each grant round robin must be unique. If the paths are considered as being a series of transitions between locations, it becomes apparent that given any starting location each round robin must make a unique transition. In order to create mutually exclusive transitions the reconnection matrix must be carefully designed. It was found that the formula below gives an ideal reconnection matrix. An example of the matrix is given in figure 2.

$$M^* = \left(M + \frac{P(P+1)}{2} \right) \bmod N$$

Where	M*	mapped pointer location
	M	round robin pointer location
	P	port number
	N	number of ports

It can be seen from the matrix in figure 2 that from any given starting position all of the pointers will move to completely different locations. Given a starting position of zero pointer five will move to location three, pointer seven will move from zero to two.

Behaviour

Certain traffic patterns can cause round robin arbiters to synchronise e.g. if one port continually requested connections to all of the other ports. Synchronisation in this particular case is not a problem, however, if there sudden fluctuation in the traffic then the rate at which the arbiter desynchronises becomes important. In these circumstances the SLIP arbiter takes N arbitration cycles to desynchronise. As the pointers in the arbiter of the invention are already reordered, the desynchronisation occurs in one arbitration cycle. The difference in performance is illustrated in figure 3.

Like a round robin, the arbiter of the invention is not inherently fair under all conditions; however, when it is used in conjunction with an arrangement which allocates weights to the requests to provide bandwidth allocation, in accordance with co-pending application No. , the combination guarantees fair allocation of connections.

Diagrams


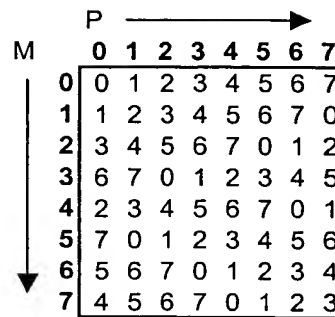


Figure 1: Example of reconnected grant signals for a four-port arbiter

THIS PAGE BLANK (USPTO)



The diagram shows an 8x8 matrix with row index M and column index P. A horizontal arrow labeled 'P' points to the right above the column headers. A vertical arrow labeled 'M' points downwards to the left of the row headers. The matrix contains the following values:

M \ P	0	1	2	3	4	5	6	7
0	0	1	2	3	4	5	6	7
1	1	2	3	4	5	6	7	0
2	3	4	5	6	7	0	1	2
3	6	7	0	1	2	3	4	5
4	2	3	4	5	6	7	0	1
5	7	0	1	2	3	4	5	6
6	5	6	7	0	1	2	3	4
7	4	5	6	7	0	1	2	3

Figure 2: Reconnection Matrix for an eight-port arbiter

THIS PAGE BLANK (USPTO)





Figure 3: Arbiter response under transitional load conditions

THIS PAGE BLANK (USPTO)